Submission to OAIC consultation on Guidance for Transparency in Automated Decision-Making

Submitted by: [redacted]

I welcome the opportunity to provide comments on the OAIC's proposed guidance for transparency in automated decision-making.

My perspective is from board, governance, legal and risk advisory work, with a particular focus on how organisations can practically govern AI, automated decision-making and privacy obligations. I am especially interested in how boards and executives can obtain confidence that automated or system-assisted decision pathways are visible, owned, explainable and capable of being challenged where appropriate.

My central submission is that the guidance should be practical, proportionate and capable of being operationalised.

The ADM transparency obligation will only be effective if organisations can apply it in a repeatable way. If the guidance creates an open-ended list of factors that must be weighed in every case, there is a risk that compliance becomes too complex, particularly for medium-sized organisations and entities that rely on third-party software platforms.

In my view, the guidance should help organisations answer five simple questions:

1. Is a computer program making, recommending, scoring, ranking, triaging or materially influencing a decision?
2. Is personal information being used in that process?
3. Could the decision reasonably be expected to significantly affect an individual's rights or interests?
4. Is the APP entity responsible for the use of the system because it built, procured, configured, enabled or relies on it?
5. Can the organisation explain the decision pathway in plain language?

That approach would support meaningful transparency without turning every automated workflow into a complex legal judgement.

Privacy policy disclosure should be treated as the endpoint of a governance process, not the starting point. Before an APP entity can make a meaningful disclosure, it needs to know where relevant automated or system-assisted decisions occur, what personal information is used, what decisions are affected, who owns the workflow, and what review or escalation pathway exists.

That points to the need for a practical ADM register. In my view, the register does not need to be overly complex. It should identify the decision type, the system used, the personal information involved, the business owner, the role played by the system, whether the decision could significantly affect rights or interests, and the relevant review or escalation pathway.

The guidance should also make clear that a "human in the loop" label is not conclusive. The practical question is whether the system output materially influences the decision in practice.

Question 1

In my view, the most important factors are:

- degree of reliance on the ADM system output: 5
- integration of ADM into the decision-making workflow: 5
- ability and likelihood of human override over an ADM decision: 5
- nature of the output, advisory vs determinative: 4
- transparency and explainability of outputs: 4

I would be cautious about adding too many further factors. The more open-ended the list becomes, the more difficult it will be for organisations to apply consistently.

The guidance should identify a small number of core indicators and then provide examples. In my view, the core indicators should be:

- whether the system output materially influences the decision;
- whether the human decision-maker has a genuine ability to depart from the system output;
- whether the system is embedded in the workflow in a way that makes reliance likely; and
- whether the organisation can explain the role played by the system in plain language.

The guidance should focus on practical influence rather than formal labels. A system may be described as "advisory", but if staff usually accept the output, or the workflow makes departure difficult, the system may still be substantially and directly related to the decision.

Question 2

Yes.

The described use of GPTea appears to be substantially and directly related to making decisions.

The chatbot is not merely supporting administration. It is summarising candidate profiles and providing recommendations about eligibility for aged-care support. Even though human staff make the final decision, the chatbot output is directly connected to the eligibility decision and appears to be a material part of the decision-making process.

Human oversight should not, by itself, remove the use case from scope. The more important question is whether the system output is relied on in practice.

In a high-impact context such as aged-care eligibility, the guidance should take a practical and cautious approach. If the system materially shapes the decision, transparency should be required.

Question 3

In my view, the listed factors should be ranked as follows:

- vulnerable persons: 5
- sensitive information: 5
- financial outcomes: 5
- intrusive practices: 5

However, I would not support the guidance becoming an extensive checklist of every possible factor that may affect rights or interests. That would make the test harder to apply and may lead to inconsistent or defensive compliance.

The guidance should instead use a simpler organising principle.

A decision is more likely to significantly affect rights or interests where it affects access to, or treatment in relation to:

- essential services;
- employment or income;
- credit, insurance, housing or financial opportunity;
- health, aged care, disability support or education;
- government benefits, licences or entitlements;
- significant personal, economic or reputational interests; or
- services or opportunities where the individual has limited ability to avoid, contest or correct the outcome.

This would give organisations a practical way to classify decisions without requiring a detailed legal analysis of every automated process.

The guidance should also recognise that significance is not limited to financial loss. A decision can significantly affect a person by restricting access, opportunity, priority, dignity, autonomy, reputation or practical ability to participate in ordinary life.

Question 4

Classes of persons who may be vulnerable should include:

- children and young people;
- elderly persons;
- persons with disability;
- persons experiencing financial hardship;
- persons with limited English proficiency;
- persons with low digital literacy;

- persons dependent on essential services, care, housing, income support or health services;
- persons with limited practical ability to opt out or choose another provider; and
- persons affected by decisions in circumstances of urgency, dependency or imbalance of power.

Again, I would favour a simple approach. The guidance should avoid turning vulnerability into a long and exhaustive list.

The key point is that vulnerability may be personal or situational. A person may be vulnerable in the context of a particular decision because they depend on the service, have limited alternatives, do not understand the process, or face serious consequences if the decision is wrong.

That practical framing would be easier for organisations to apply than a rigid list of categories.

Question 5

In my view, the listed services and supports should be ranked as follows:

- physical assistance/support: 5
- financial assistance: 5
- access to educational services: 5
- in-home services: 5
- access to essential banking and credit services: 5
- access to telecommunications: 5
- access to essential utilities: 5

The guidance should focus on whether the service or support is significant to the individual, rather than only on the type of service.

A service may be significant because it is essential, because alternatives are limited, because delay causes harm, or because the person is dependent on the service at the relevant time.

Examples should include healthcare, aged care, disability support, housing, insurance, employment, government benefits, utilities, telecommunications, credit and essential digital access.

However, the guidance should avoid creating an overly detailed taxonomy. The more useful approach would be to identify broad categories of significant services and then provide illustrative examples.

Question 6

The OAIC should have regard to a range of legal frameworks and policies when considering rights and interests, including:

- anti-discrimination law;
- administrative law principles, including reviewability and procedural fairness;
- employment law;
- work health and safety law;
- consumer protection law;
- credit, banking and insurance regulation;
- human rights legislation in relevant jurisdictions;
- child protection, aged care, disability and health regulation;
- freedom of information principles for government agencies;
- procurement and outsourcing obligations for government and regulated entities;
- responsible AI policies and standards; and
- guidance from the Commonwealth Ombudsman on automated decision-making.

However, the guidance should not require organisations to conduct a full legal review across all potentially relevant regimes each time they classify a system.

The APP 1 transparency obligation should be framed as a practical privacy transparency obligation. It sits alongside other obligations, but it should not become a proxy for every other legal or regulatory issue.

For boards and executives, the practical question should be whether the organisation can identify, explain and govern the relevant decision pathways. Where a decision pathway touches higher-risk areas, further legal or assurance review may be appropriate.

Question 7

Yes, differential pricing for a consumer product may significantly affect an individual's interests, depending on the circumstances.

I would not suggest that every instance of price variation should be treated as significant. However, in the edge case, the difference between $100 and $35 is substantial. It is based on an inferred proxy for wealth, using the individual's address, and it has the practical effect of excluding Marco from purchasing the product.

I would be cautious about adopting a single fixed percentage threshold. A percentage threshold may be useful as a practical indicator, but it should not be the only test.

The more useful factors are:

- whether the product or service is important or essential;
- whether the price difference is material in dollar terms;
- whether the individual has realistic alternatives;

- whether the pricing is based on sensitive information or proxies for sensitive attributes;
- whether the pricing practice is reasonably expected; and
- whether the result is exclusionary or unfair in practice.

For the factors listed:

- dollar amount of differential pricing: 4
- significance of the product or service being purchased: 5

The significance of the product or service should generally carry greater weight than the percentage difference alone. Differential pricing for healthcare, insurance, credit, housing, utilities or essential digital access should be treated differently from differential pricing for a discretionary product.

Question 8

Yes.

It should be considered a decision that Rafqa did not receive the Manscon job advertisement.

The relevant decision is not only the final hiring decision. The algorithmic promotion of the job advertisement determines who is given a meaningful opportunity to apply. If a person is excluded from seeing the opportunity because of an automated targeting process, that exclusion can materially affect employment interests.

This is particularly important where the targeting uses gender, or a proxy for gender, in a way that reinforces existing workforce imbalance.

The guidance should recognise that ADM can affect rights or interests by controlling visibility, ranking, prioritisation, triage, eligibility or access, even before a formal final decision is made.

Question 9

Yes. The meaning of “arranged for” requires practical guidance, particularly in relation to third-party systems and embedded functionality.

Many organisations will not build ADM systems themselves. They will procure, configure or use systems supplied by third parties. This is one of the most important practical issues for boards and executives.

The guidance should make clear that “arranged for” should not be limited to systems designed or owned by the APP entity. If an entity procures, deploys, configures, enables or relies on a computer program in a decision pathway, it may have arranged for that program to make, or do something substantially and directly related to making, a decision.

Examples may include:

- SaaS platforms with embedded ADM functionality;
- recruitment platforms that rank, screen or target candidates;

- insurance, banking or credit platforms that score or triage applicants;
- customer service platforms that classify, prioritise or route matters;
- fraud detection, identity verification or risk-scoring systems;
- procurement platforms that score suppliers or bids;
- HR systems that make recommendations about employees or applicants; and
- software updates that introduce ADM functionality after the original procurement decision.

At the same time, the guidance should be practical. Organisations often have limited visibility into proprietary systems. The obligation should not assume that APP entities have perfect technical knowledge of every vendor model.

In those cases, the guidance should encourage proportionate governance, including vendor due diligence, contractual controls, vendor attestations, records of configuration decisions, risk assessment and plain-language disclosure based on what the organisation knows or ought reasonably to know.

Question 10

Yes. The guidance should encourage disclosures that are meaningful without being overwhelming.

In my view, privacy policy disclosures should explain:

- the kinds of decisions affected by ADM;
- the kinds of personal information used;
- whether the computer program makes the decision or materially assists a human decision-maker;
- the practical role of the system, such as scoring, ranking, recommending, triaging, prioritising or flagging;
- whether third-party systems or service providers are involved;
- how an individual can seek further information; and
- how an individual can request review or raise a concern, where available.

The guidance should discourage generic statements such as “we may use automated technologies to improve our services”. That type of wording is unlikely to provide meaningful transparency.

At the same time, disclosure should not become so detailed that it becomes unreadable. Privacy policies are already difficult for many individuals to understand. The objective should be plain-language transparency about the kinds of decisions and personal information involved, not technical explanation of every system or algorithm.

A useful disclosure should allow an individual to understand:

- whether ADM may affect them;
- what kind of decision is involved;
- what kind of personal information is used;
- whether ADM is determinative or influential;
- what they can do if they are concerned; and
- where they can go for review or further information.

For boards and executives, the practical question should not be only whether the privacy policy has been updated. It should be whether the organisation can identify, explain and evidence the automated decision pathways sitting underneath the disclosure.

Concluding comments

The ADM transparency obligation is important, but it needs to be capable of practical implementation.

The guidance should not create an open-ended compliance exercise that requires organisations to weigh an expanding list of factors for every system. That would make compliance harder, not better.

The better approach is a simple and repeatable triage model:

1. identify automated or system-assisted decision pathways;
2. determine whether personal information is used;
3. assess whether the system materially influences the decision;
4. assess whether the decision could significantly affect rights or interests;
5. identify who owns the workflow and review pathway; and
6. disclose the relevant classes of decisions and personal information in plain language.

This approach would support meaningful transparency, improve organisational accountability and give boards and executives a practical way to oversee ADM risk.